







BDDs Strike Back

Efficient Analysis of Static and Dynamic Fault Trees

Daniel Basgöze¹, Matthias Volk²(✉) , Joost-Pieter Katoen^{1,2} ,
Shahid Khan¹ , and Marielle Stoelinga^{2,3} 

¹ Software Modeling and Verification, RWTH Aachen University,
Aachen, Germany

² Formal Methods and Tools, University of Twente, Enschede, The Netherlands
m.volk@utwente.nl

³ Department of Software Science, Radboud University, Nijmegen, The Netherlands

Abstract. Fault trees are a key model in reliability analysis. Classical static fault trees (SFT) can best be analysed using binary decision diagrams (BDD). State-based techniques are favorable for the more expressive dynamic fault trees (DFT). This paper combines the best of both worlds by following Dugan’s approach: dynamic sub-trees are analysed via model checking Markov models and replaced by basic events capturing the obtained failure probabilities. The resulting SFT is then analysed via BDDs. We implemented this approach in the STORM model checker. Extensive experiments (a) compare our pure BDD-based analysis of SFTs to various existing SFT analysis tools, (b) indicate the benefits of our efficient calculations for multiple time points and the assessment of the mean-time-to-failure, and (c) show that our implementation of Dugan’s approach significantly outperforms pure Markovian analysis of DFTs. Our implementation STORM-DFT is currently the only tool supporting efficient analysis for both SFTs and DFTs.

1 Introduction

Fault trees [47, 51, 54] are a common formalism in reliability engineering and required by standards in a broad range of industries [27, 33, 51]. A fault tree represents a Boolean function and models how overall system failures depend on the failure of basic system components. Fault tree analysis (FTA) is commonly performed by translating a fault tree into a binary decision diagram (BDD) and calculating the relevant metrics on this BDD [43, 50]. BDDs yield compact representations of fault trees enabling the analysis of large systems [17]. Ongoing improvements in BDD tools such as parallelisation [20] allow for modern implementations of FTA via BDDs.

This work has been partially funded by NWO under the grant [PrimaVera](#) number NWA.1160.18.238, European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101008233 (*Mission*), and the ERC Consolidator Grant 864075 (*CAESAR*). Khan is funded by a HEC-DAAD stipend.

© Springer Nature Switzerland AG 2022

J. V. Deshmukh et al. (Eds.): NFM 2022, LNCS 13260, pp. 713–732, 2022.

https://doi.org/10.1007/978-3-031-06773-0_38

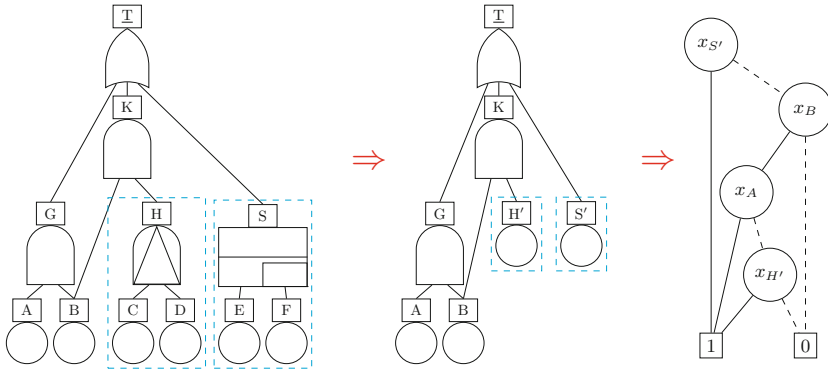


Fig. 1. DFT modularisation

Fault trees are static in nature and their expressiveness is limited. *Dynamic fault trees (DFT)* [21] support ordered failures, spare management and functional dependencies. The flexibility and increased expressiveness of DFTs however requires more involved analysis methods. BDDs cannot represent DFTs directly as DFTs consider failure sequences instead of Boolean combinations. A common approach is to translate DFTs into Markov models [10, 22, 56]. Gulati and Dugan [31] proposed a modular DFT analysis approach combining several analysis techniques. It divides the DFT into independent sub-parts which are analysed individually. *Modularisation thus allows to use the best of both worlds: Markov models for dynamic parts and BDDs for static parts.*

The idea of modularisation is depicted in Fig. 1. First, dynamic modules, i.e., sub-trees containing dynamic elements, are identified (the blue boxes in the left-most tree). Each dynamic module is analysed independently with state-of-the-art analysis techniques for Markov models [56]. Afterwards, each dynamic module is replaced by a single basic event which represents the corresponding failure probabilities (tree in the middle). The remaining (static) fault tree is then translated into a BDD (right most part) and is analysed by BDD techniques. Modularisation works especially well when the dynamic parts are contained at the bottom of the fault tree and the static parts are on top. As dynamic parts commonly model single components, this structure is present in most DFTs.

Related Work. FTA via BDDs was first presented in [43] and [17], and successively improved in [25, 50]. BDD-based analysis of static fault trees (SFTs) is supported by academic tools such as SCRAM [40], XFSTA [45] and SHARPE [55], as well as commercial tools, e.g., RISK SPECTRUM [3]. We refer to [47] for a detailed overview on BDD-based FTA.

For DFTs, various analysis techniques exist. Common approaches translate DFTs into models such as Markov models [10, 22, 56], Bayesian networks [38] and Petri nets [39], or are based on Monte-Carlo simulation [15, 41]. There also exist analysis techniques based on extensions of BDDs such as sequence decision diagrams [44], sequential BDDs [58], multiple-valued decision diagrams [36]

or conditional BDDs [60]. These approaches use BDDs to enumerate all failure sequences leading to a system failure so as to compute the overall system unreliability. To the best of our knowledge, no tool support exists for any of the BDD approaches for DFTs and their scalability for large DFTs remains unclear.

For SD fault trees [35]—another extension of static fault trees with limited dynamic behaviour—efficient analysis techniques exist via minimal cut sets [35] and abstraction [4]. However, expressiveness of SD fault trees is limited compared to DFTs as they only allow dynamic behaviour in basic system components.

Our DFT approach is based on modularisation [31] which was first implemented in the DIFTREE tool [23] and its successor GALILEO [53]. However, both tools are not available anymore for more than a decade. Recent work [56] has implemented modularisation for DFTs but is limited to independent sub-trees that must be direct children of the top event. In addition, [56] analyses static FTs by translation to Markov models.

Implementation. We implemented the BDD translation for static fault trees (SFTs) in the STORM model checker [32] and use the multi-core BDD library SYLVAN [20]. Our implementation STORM-DFT supports computing minimal cut sets (MCS), the unreliability and several importance measures such as the Birnbaum index [8]. STORM-DFT exploits vectorisation and thus enables to compute a metric for multiple time bounds at once. In addition, we support the calculation of the mean-time-to-failure (MTTF) via approximation. For DFTs, we implemented the modularisation approach exploiting both the BDD translation and our efficient DFT analysis via Markov models [56].

Evaluation. Experiments on a benchmark set of 215 SFTs and 124 DFTs yield:

- On SFTs, STORM-DFT is competitive compared to existing tools such as SCRAM [40] and XFTA [45] and is significantly faster when analysing multiple time points due to vectorisation.
- The variable ordering in STORM-DFT is not yet optimal and can lead to larger BDDs and subsequently longer run times.
- On DFTs, our BDD-based modularisation is significantly faster than both plain Markov-model analysis [56] and an existing realisation of (top-down) modularisation [56] based on pure Markov-model analysis.

Our implementation is publicly available in the open-source tool STORM-DFT¹. We also provide an artifact of the experimental evaluation containing the scripts, tool configurations and fault tree models².

Contributions. In summary, the main contributions of this paper are:

- A competitive implementation of SFT analysis via BDDs in STORM [32].
- Vectorisation for multiple time bounds and an approximation for MTTF.
- A fast implementation of a modern version of modularisation using BDDs.

Our implementation STORM-DFT is the only state-of-the-art analysis tool for both static and dynamic fault trees.

¹ <https://www.stormchecker.org/>.

² <https://doi.org/10.5281/zenodo.6390998>.



Fig. 2. Node types in ((a)–(d)) static and (all) dynamic fault trees.

Structure of the Paper. We introduce fault trees and binary decision diagrams in Sect. 2. Section 3 presents the analysis of (static) fault trees using BDDs. We evaluate the approach in Sect. 4. Section 5 presents the analysis of dynamic fault trees via modularisation and using BDDs for static parts. We evaluate the approach in Sect. 6. We conclude in Sect. 7 and present future work.

2 Preliminaries

2.1 Fault Trees

Fault trees (FTs) model how failures can occur and propagate in systems [47, 51, 54]. FTs are *directed acyclic graphs* in which the leaves are called *basic events* (BEs) and intermediate nodes are called *gates*. A BE represents an (atomic) system component which fails according to a given failure distribution. Failures of BEs are propagated through the system according to the gates and eventually lead to a failure of the unique root of the graph, the *top event*. *Dynamic fault trees (DFTs)* [21] are the most prominent extension of fault trees and their additional gates allow for more realistic modelling. Note that we do not consider repairs.

Definition 1 (Dynamic fault tree). A dynamic fault tree (DFT) is a tuple $\mathcal{F} = (V, \sigma, \text{Type}, \text{top}, \Theta)$ where

- V is a finite set of nodes.
- $\sigma : V \rightarrow V^*$ defines the ordered children of a node (also called the inputs).
- $\text{Type} : V \rightarrow \{\text{BE}, \text{AND}, \text{OR}, \text{VOT}_k\} \cup \{\text{PAND}, \text{POR}, \text{PDEP}, \text{SEQ}, \text{SPARE}\}$ defines the type of a node.
- $\text{top} \in V$ is the top event.
- $\Theta : \{v \in V \mid \text{Type}(v) = \text{BE}\} \rightarrow \Omega$ maps each BE to a failure distribution from Ω the set of probability distributions.

A VOT_k -gate satisfies $1 \leq k \leq |\sigma|$. A *static fault tree (SFT)* is a DFT where the node types Type are restricted to $\{\text{BE}, \text{AND}, \text{OR}, \text{VOT}_k\}$. In DFTs, we restrict the failure distributions of BEs to exponential distributions to allow for analysis based on Markov models. We say a “DFT \mathcal{F} is failed” if top is failed.

We shortly introduce the different node types; the precise semantics is given in [34]. The graphical representation of each node type is given in Fig. 2.

Basic events (BEs) fail according to their associated failure distributions. Commonly, an exponential failure distribution with a failure rate λ is used.

Static gates represent Boolean logic functions. The AND-gate fails if all its inputs fail. The OR-gate fails if at least one input fails. The VOT_k -gate is a generalisation and fails if at least k inputs fail.

Priority gates extend the static gates with the additional constraint that the inputs have to fail in order from left to right. Failures out of this order render the gate *fail-safe* and it can never fail. The PAND-gate fails if all inputs fail from left to right. The POR-gate fails if the leftmost child fails before all other gates.

Dependencies encode functional dependencies of the system. If the first child of the PDEP_p fails, all other children fail with probability p .

Sequence enforcers ensure that children only fail in order from left to right.

Spare gates model spare management. Initially, the first child is used. If it fails, the next child is claimed and used, and so forth. The SPARE fails if all its children failed. Children can be shared by multiple SPAREs, but can only be used exclusively by one SPARE. Using a child activates the corresponding components and can increase the associated failure rate.

Example 1 (Fault tree). Consider the DFT on the left of Fig. 1. The top event T fails for example if both BEs A and B fail. The DFT also fails if B and PAND H fail. The PAND only fails if the first child C fails before D .

Fault Tree Analysis. Fault trees are analysed w.r.t. the failure of the top event. Common metrics are the *unreliability* within a given time bound and the *mean-time-to-failure (MTTF)*. For SFTs, the *minimal cut sets* play an important role. Cut sets with only a few elements for example indicate system vulnerabilities.

Definition 2 (Minimal cut sets). *Let \mathcal{F} be a static fault tree. A minimal cut set (MCS) for \mathcal{F} is a set $M \subseteq \text{BE}$ such that:*

1. *the failure of all BEs in M leads to the failure of \mathcal{F} , and*
2. *M is minimal, i.e., no subset $M' \subsetneq M$ leads to the failure of \mathcal{F} .*

2.2 Binary Decision Diagrams

Binary decision diagrams (BDDs) [1, 13] are graphs based on the Shannon expansion [48]. We introduce BDDs by example and refer to [14] for more details.

A BDD \mathcal{B} encodes a Boolean function f over variables x_1, \dots, x_n . Nodes in \mathcal{B} represent variables of f and follow a given variable ordering. Outgoing edges of a node x represent the two possible assignments of variable x : the solid line represents $x = 1$, the dashed line $x = 0$. Leaves represent functions 1 and 0.

Example 2 (BDD). Consider the BDD on the right of Fig. 1. The BDD represents the function $f = x_{S'} \vee (x_B \wedge (x_A \vee x_{H'}))$. The satisfying assignments of f are obtained by following all paths from the root to the 1-leaf.

3 SFT Analysis via BDD

Translation from SFT into BDD. An SFT \mathcal{F} is translated into a BDD by simply calculating the BDD-representation of the propositional formula representing the

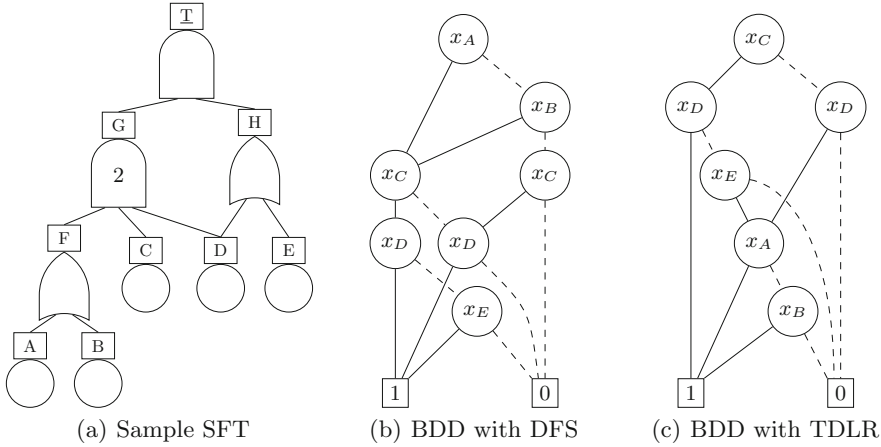


Fig. 3. SFT and corresponding BDDs for different variable orders

failure behaviour of \mathcal{F} [43]. The algorithm follows a simple recursive bottom-up approach which combines the BDDs of sub-trees according to the logic gates. The VOT_k -gate is translated by exploiting the Shannon decomposition [48].

Variable Ordering. The variable ordering employed for the BDD is important as different orderings can result in significantly different BDD sizes [2, 49]. Common variable orderings are *depth-first search (DFS)* [43] and *top-down left-right (TDLR)* [2]. Finding optimal variable orderings is still ongoing research [11, 42].

Example 3 (BDDs for SFT). Consider the SFT depicted in Fig. 3(a). The corresponding BDD for the DFS ordering $x_A <_{\text{Var}} x_B <_{\text{Var}} x_C <_{\text{Var}} x_D <_{\text{Var}} x_E$ has 7 nodes and is given in Fig. 3(b). The BDD for variable ordering TDLR $x_C <_{\text{Var}} x_D <_{\text{Var}} x_E <_{\text{Var}} x_A <_{\text{Var}} x_B$ has 6 nodes and is depicted in Fig. 3(c).

3.1 Computing Minimal Cut Sets

Minimal cut sets are a common metric for SFTs. Several approaches exist to compute MCS [47], and we focus on the BDD-based approach [43]. We are interested in all paths of the BDD that reach the 1-leaf, i.e., lead to a failure of the SFT. All variables reached by the 1-edge on such a path form a *solution* of the BDD. Each solution is a cut set. The aim is to compute all minimal solutions, i.e., solutions whose proper subsets are not a solution.

A naïve approach computes the solutions for each (sub-)BDD in a bottom-up way. The solutions of a node v are then the union of the solutions of the 1-successor of v extended with v , and the solutions of the 0-successor. However, the resulting solutions are not necessarily minimal [6, 43].

The algorithm of [43] exploits the fact that SFTs can only encode monotonic switching functions (a.k.a. *coherent* FTs). From the original BDD \mathcal{B} , a new BDD \mathcal{B}' is constructed whose solutions are exactly the minimal solutions of \mathcal{B} .

This construction uses the ‘without’-operator on BDDs to exclude parts of the 1-successor which are already included in the 0-successor; see [6, 43] for details.

3.2 Computing Unreliability

Let X be the random variable representing the failure of BE x . We use $P_t(f_x) := P(X \leq t)$ to denote the probability that x fails within time bound t . A common metric is the *unreliability* $P_t(\mathcal{F}) := P_t(f_{top})$, i.e., the probability that the SFT \mathcal{F} with TLE top fails within time bound t . This metric can easily be computed on the BDD by employing Shannon decomposition: $P_t(f) = P_t(f_x) \cdot P_t(f|_{x=1}) + (1 - P_t(f_x)) \cdot P_t(f|_{x=0})$. The algorithm works independently of the calculation for $P_t(f_x)$ and can therefore be applied to any failure distribution of the BEs.

Computing Sensitivity Measures. Importance measures [47, 54] are used to assess how sensitive the overall system is w.r.t. sub-systems. The *Birnbaum importance index* [8] is a prominent metric. For BE e in SFT \mathcal{F} at time t , the Birnbaum index is given by the conditional probabilities $BI_t(\mathcal{F}, e) := P_t(\mathcal{F} | e) - P_t(\mathcal{F} | \neg e)$, where $\neg e$ represents that e has not failed. The calculation is done on the BDD corresponding to \mathcal{F} , cf. [25]. Additional metrics are computed based on the Birnbaum index [25], e.g., the *critical importance factor*, the *Vesely-Fussell importance factor* [28], the *risk achievement worth* and the *risk reduction worth* [16].

Vectorisation for Multiple Time Bounds. Investigating the unreliability over time requires computing the unreliability for a large number of different time bounds. The calculation for multiple time bounds can easily be parallelised, because the computations are independent: the failure probability of the top event within time bound t only depends on the failure probabilities of the BEs within t . We employ *vectorisation* [26] where multiple probabilities—corresponding to different time bounds—are stored within a single vector and computed concurrently. Vectorisation exploits both temporal and spatial locality in modern CPU caches as well as *SIMD instructions* (single instruction multiple data) which operate on arrays of values at once.

3.3 Computing the MTTF

Vectorisation naturally leads to approximation methods as we can efficiently evaluate many time bounds at once. We exploit this for the *mean-time-to-failure (MTTF)*, i.e., the expected time point the SFT \mathcal{F} fails at. It is calculated by $\int_0^\infty P_t(\mathcal{F}) dt$. We numerically approximate the improper integral by sampling a large number of time bounds on the BDD obtained for the SFT. We use two different methods from [18]. The first method, *proceeding to the limit*, computes a sequence of integrals $\int_{r_i}^{r_{i+1}} P_t(\mathcal{F}) dt$ until the result of an integral is less than a given error ε . Our implementation uses varying steps sizes which start at 10^{-10} and a default error of $\varepsilon := 10^{-12}$. The second method, *change of variable*, aims to “squeeze” the unbounded interval $[0, \infty)$ into the bounded interval $[0, 1)$ using integration by substitution. The latter method always uses the same number of samples (default 10^6) and works good for functions slowly approaching zero.

The former method uses a variable number of samples and performs better for functions which approach zero relatively fast or change rapidly; see [6] for details.

3.4 Implementation

We implemented the SFT analysis in the STORM-DFT³ tool based on the STORM model checker [32] and use the multi-core library SYLVAN [20] for creating and handling BDDs. We list the main implementation details in the following.

Multi-core BDD. SYLVAN natively enables multi-core computations on BDD [19]. STORM-DFT exploits this when performing the translation from SFT to BDD.

Complement Edges. The implementation uses *complement edges* [12] which negate the corresponding function and it allows to use a single terminal node.

Variable Ordering. The implementation uses the order of BEs given in the input file (in Galileo format⁴) as the variable ordering for the BDD. That way, we support arbitrary variables orderings which can be explicitly given by either the user or a pre-processing step. Currently, STORM-DFT supports the DFS and TDLR variable orderings via pre-processing steps.

Caching. During the translation, intermediate BDDs are not cached in order to reduce the memory consumption. Caching can be explicitly enabled for specific events if needed.

Vectorisation. STORM-DFT uses the EIGEN library [30] for vectorisation. The *chunk-size*, i.e., the number of time points computed in parallel, can be configured from the command-line. We refer to [6] for details on the optimal chunk-size.

Properties. Apart from standard metrics presented before, our implementation supports properties defined in a fragment of *continuous stochastic logic (CSL)* [5]. More precise, STORM-DFT supports (time-bounded) reachability formulas of the form $\mathbb{P}_{=?}(\diamond^{\leq t}\phi)$ for state formula ϕ and time bound t .

4 Evaluation of SFT Approach

We evaluate the fault tree analysis via BDDs as implemented in STORM-DFT on a range of benchmarks and compare with existing tools. For reproducibility, we provide an artifact online⁵ which contains the analysis scripts, tool configurations and fault tree models used in our experimental evaluation.

4.1 Configurations

We use STORM-DFT version 1.6.3. In the default configuration, STORM-DFT is single-threaded, uses a chunk-size of 1024 for vectorisation and uses the DFS variable ordering. We also evaluate STORM-DFT in a configuration using the TDLR variable ordering and in configurations using multiple cores.

³ <https://www.stormchecker.org/>.

⁴ <https://dftbenchmarks.utwente.nl/galileo.html>.

⁵ <https://doi.org/10.5281/zenodo.6390998>.

Table 1. SFT benchmark sizes

	Aralia	Sprinkler	Railway	Industry	Random	Random (large)
#BEs	25–1567	31	22–54	36–184	150	500
#Gates	20–1622	35	69–259	21–67	70–122	261–316

We compare the SFT analysis in STORM-DFT with two existing tools: SCRAM⁶ (version 0.16.2) and XFATA⁷ (version 2.0.1).

SCRAM [40] is an open-source probabilistic risk analysis tool which supports the *Open-PSA Model Exchange* format [52]. It performs SFT analysis using BDDs and supports metrics such as MCS, unreliability and importance measures. Before analysis, SCRAM simplifies the SFT’s graph structure. The BDD variable ordering then follows the topological ordering on the simplified SFT.

XFATA [45] is a free-to-use tool for the analysis of fault trees and similar models. It is hosted by the AltaRica Association. XFATA uses its own object-oriented design language *S2ML+SBE* as input, but also supports the Open-PSA format. The analysis is performed by either generating the MCS and calculating the metrics on them or by creating a BDD and computing the metrics via the BDD. For the variable ordering, the children of gates are sorted beforehand and then the DFS ordering is used.

4.2 Benchmarks

We use the following collection of SFT benchmarks for our evaluation:

- 40 examples from the Aralia benchmark set⁸. We excluded 3 non-coherent SFTs containing a negation-gate as STORM-DFT does not support them.
- 3 models of wet-pipe fire sprinkler systems in Australian shopping centres [37]
- 8 examples modelling train routing options w.r.t. infrastructure failures in railway station areas [57].
- 3 industrial models for components of a lock used in water navigation.
- 161 randomly generated SFTs using a script provided by the SCRAM tool. 128 of the random SFTs have 150 BEs and 33 are large SFTs with 500 BEs.

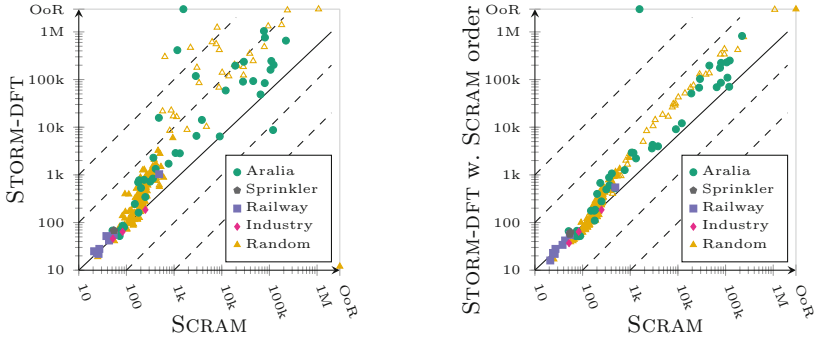
We provide statistics on the benchmarks in Table 1. We give the minimal and maximal number of BEs and gates for each benchmark set.

We analyse each fault tree w.r.t. four queries: all minimal cut sets (*MCS*), the *unreliability* at time point $t = 1$, the *unreliability for 10 000 time points* that are uniformly distributed within the interval $[0, 10]$, and the *Birnbaum importance index* at time point $t = 1$.

⁶ <https://github.com/rakhimov/scram>.

⁷ <https://altarica-association.org/members/arauzy/Software/XFATA/XFATA2.html>.

⁸ <https://github.com/rakhimov/scram/tree/develop/input/Aralia>.



(a) BDD sizes STORM-DFT DFS vs SCRAM (b) BDD sizes STORM-DFT using variable ordering from SCRAM vs SCRAM

Fig. 4. Comparison of BDD sizes for different variable orderings for MCS

4.3 Results

We ran STORM-DFT, SCRAM and XFTA on all 215 examples w.r.t. the four different queries. We ran the experiments on a desktop machine with an AMD Ryzen™ 9 5950X and 32 GB of RAM running Arch Linux. In the multi-core configuration, we used 16 cores. The timeout was set to 5 min and the memory was limited to 30 GB. We asserted that the obtained results are the same for all three tools. In the following, we provide detailed comparisons of the tools. Additional results and details can be found in the extended version of this paper [7].

We present the comparisons as scatter plots such as in Fig. 5. All scatter plots are in log-log scale and indicate—in most cases—the time (in seconds) it took each tool to compute a query. Line *OoR* indicates *out of resources* and represents either a timeout or memory out. All points below the diagonal indicate examples which STORM-DFT could solve faster than the other tool. All points below the first (second) dashed line correspond to SFTs for which STORM-DFT was one (two) order(s) of magnitude faster than the other tool. Similarly, for every point above the diagonal, the other tool was faster.

BDD Sizes. First, Fig. 4 compares the number of nodes in the BDDs which provides an idea of the respective memory consumption. Figure 4(a) compares the sizes of the BDDs obtained by STORM-DFT and SCRAM. The largest BDDs which could be analysed contain more than a million nodes. In general, SCRAM yields smaller BDDs, for larger BDDs even by more than one order of magnitude. The main reason is that SCRAM uses a slightly different variable ordering which seems to yield smaller BDDs.

The influence of the variable ordering is further investigated in Fig. 4(b). Here, we extract the variable ordering from SCRAM and employ it in STORM-DFT. We see that using the SCRAM variable ordering in STORM-DFT improves upon the default DFS ordering and yields smaller BDD sizes—in particular for larger SFTs. However, SCRAM still yields smaller BDDs than STORM-DFT.

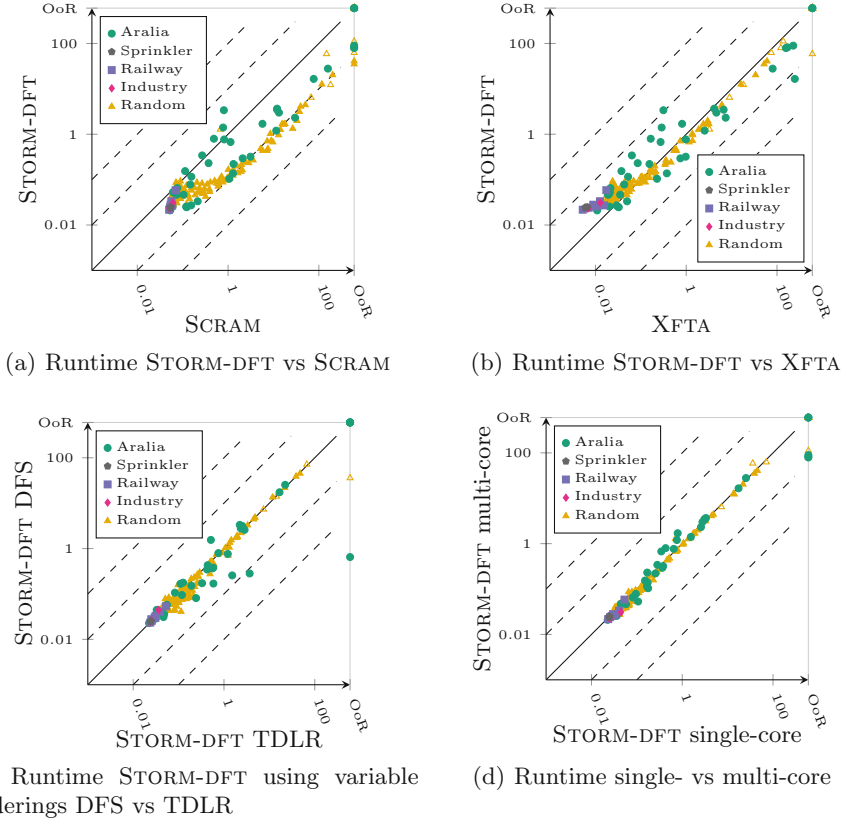


Fig. 5. Comparisons for the computation of MCS

Reasons for the discrepancy could be that the BDD implementation in SCRAM is specifically tailored to SFTs or that the self-reported number of BDD nodes in both tools are computed in different ways.

Using good heuristics for the variable ordering is crucial for a small memory footprint; the heuristics in STORM-DFT can be further improved.

MCS. We compare the runtimes for computing the MCS in Fig. 5. Figure 5(a) compares STORM-DFT and SCRAM. We first see that the SFTs corresponding to the sprinkler, railway and industry case studies can be solved within 1 s by both tools. This also holds for other tools/configurations and metrics. As these SFTs are not a challenge, we focus on the Aralia benchmark and random SFTs in the remainder. STORM-DFT is faster than SCRAM in nearly all cases. One possible reason is that SCRAM outputs the MCS in an XML format which requires more I/O-operations than the simple list output of STORM-DFT.

When comparing STORM-DFT with XFTA (cf. Fig. 5(b)), the picture is more diverse. XFTA is faster than STORM-DFT on most examples which can be solved within 1 s. This is mostly due to the overhead resulting from initializing the

SYLVAN BDD library within STORM. For some of the medium-sized Aralia examples, XFTA performs better than STORM-DFT. However, for larger examples, STORM-DFT prevails on all the random SFTs and nearly all Aralia benchmarks. *For MCS, STORM-DFT is faster than both SCRAM and XFTA for larger SFTs.*

Figure 5(c) compares the runtimes of STORM-DFT for different variable orderings DFS and TDLR. While both variable orderings do not make much of a difference for most examples, DFS performs better for some of the fault trees and even allows to handle an FT which is OoR for TDLR.

Figure 5(d) shows that using 16 cores (instead of a single core) for the BDD operations as supported by the SYLVAN library has only a minor influence on the runtime. One reason is that most operations performed on the BDDs are fairly basic and therefore do not profit much from parallelization. However, for some large examples, the configuration with multiple cores allows to handle SFTs which were OoR for the single core.

Unreliability. Figure 6 shows the runtimes for computing the unreliability. Figure 6(a) compares STORM-DFT and SCRAM. For most examples, both tools compute the unreliability within 1s. For larger benchmarks, SCRAM outperforms STORM-DFT. The main reason is that STORM-DFT builds larger BDDs than SCRAM, cf Fig. 4(a). XFTA is faster than STORM-DFT on the small examples, cf. Fig. 6(b). However, for larger SFTs, STORM-DFT outperforms XFTA. For all tools, computing the unreliability is significantly faster than computing the MCS.

Figure 6(c) compares the performance of STORM-DFT and SCRAM when computing the unreliability for 10 000 different time points. STORM-DFT performs vectorisation with a chunk-size of 1024 and thus computes 1024 time points at once. This dedicated support yields a clear performance gain compared to SCRAM which computes each time point sequentially. For larger SFTs, STORM-DFT is more than one order of magnitude faster. The same holds true when comparing to XFTA, cf. Fig. 6(d).

STORM-DFT is slower when computing the unreliability for one time bound, but is significantly faster than SCRAM and XFTA for multiple time bounds.

Importance Measures. Last, we consider the computation of the Birnbaum importance index for all BEs in a SFT. We omit the results for SCRAM as SCRAM needs to compute the MCS for computing the Birnbaum importance index for all BEs. As STORM-DFT does not need this computation, the comparison would be unfair. We provide the results in the extended version of this paper [7].

Figure 7(a) compares the runtime for STORM-DFT and XFTA when computing the Birnbaum importance index for all BEs at a single time point. We see that most examples are solved within 1s. XFTA performs better on the larger examples. Figure 7(b) compares both tools when computing the Birnbaum importance index for 1000 time points. Here, STORM-DFT is orders of magnitude faster than XFTA and provides results where XFTA runs out of resources.

When computing the Birnbaum importance index for all BEs, XFTA is faster for single time points whereas STORM-DFT is orders of magnitude faster for multiple time points.

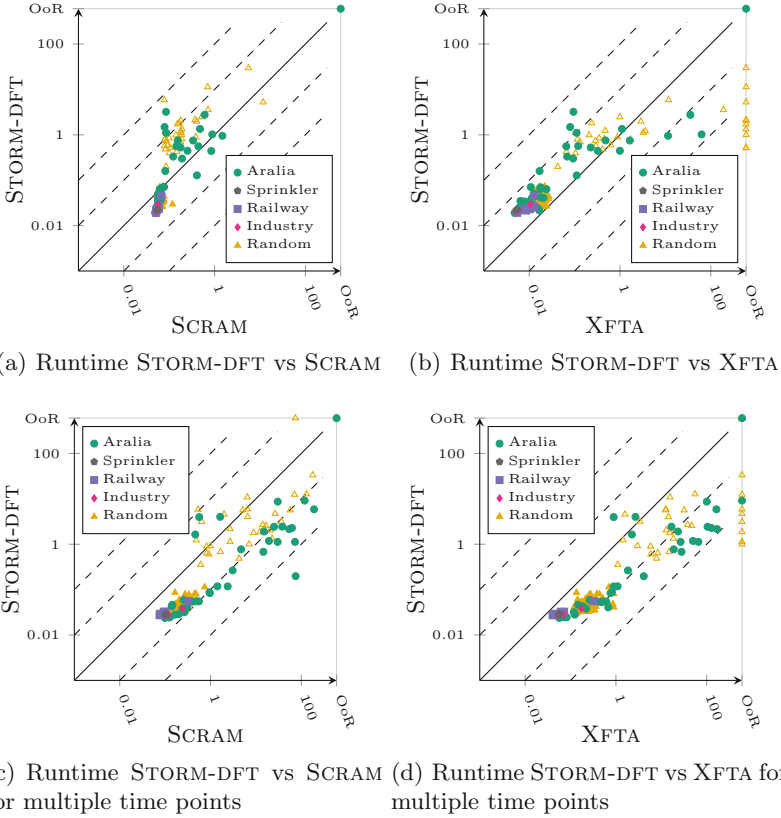


Fig. 6. Comparisons for the computation of the unreliability

5 DFT Analysis via BDD and Modularisation

Dynamic fault trees extend SFTs by capturing dynamic failure behaviour such as ordered failures, spare management or functional dependencies. Analysis of DFTs therefore needs to keep track of the history of failures and BDDs cannot be easily used. In this approach, we combine SFT and DFT analysis using *modularisation* [31]. Modularisation is a “divide-and-conquer”-approach which splits the DFT into *modules*, i.e., independent sub-trees. Each module is analysed independently and the corresponding results are combined in the end. Modularisation thus allows to exploit the “best” analysis technique for each module individually.

Modules containing dynamic elements are analysed by translating the corresponding sub-DFT into a Markov model [56]. The state space of the Markov model is created by exhaustively exploring all possible BE failures of the DFT. Each transition corresponds to the failure of a BE and the successor state represents the status of the DFT after the BE failure. The transition rate is given by the failure rate of the BE. Our translation employs several optimisation tech-

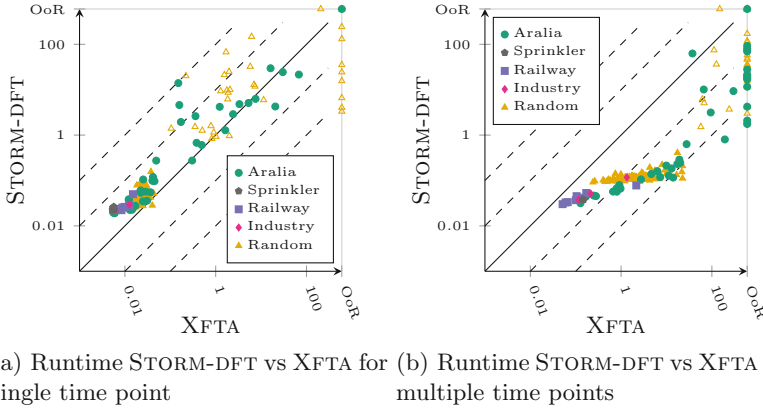


Fig. 7. Comparison for computing the Birnbaum importance index for all BEs

Algorithm 1. DFT analysis via modularisation

Input: DFT \mathcal{F} , time bounds t_1, \dots, t_n

Output: Analysis results $P_{t_1}(\mathcal{F}), \dots, P_{t_n}(\mathcal{F})$

Compute the modules $D = \{\mathcal{F}_1, \dots, \mathcal{F}_k\}$ in \mathcal{F}

for $\mathcal{F}_i \in D$ sorted by decreasing size of \mathcal{F}_i **do**

if $\mathcal{F}_i \setminus \bigcup_{\mathcal{F}' \neq \mathcal{F}_i} \mathcal{F}'$ contains no dynamic gate **then**
 $D := D \setminus \mathcal{F}_i$

for $\mathcal{F}_i \in D$ **do**

Generate Markov model \mathcal{C}_i from \mathcal{F}_i

Compute failure probabilities $p_1 = P_{t_1}(\mathcal{C}_i), \dots, p_n = P_{t_n}(\mathcal{C}_i)$ on \mathcal{C}_i

Create BE B_i such that $P_{t_j}(B_i) = p_j$ for all $1 \leq j \leq n$

Replace \mathcal{F}_i by B_i in \mathcal{F}

Build BDD \mathcal{B} from \mathcal{F}

Compute results $r_1 = P_{t_1}(\mathcal{B}), \dots, r_n = P_{t_n}(\mathcal{B})$ on \mathcal{B}

return r_1, \dots, r_n

niques to mitigate a state space explosion. The optimisations encompass discarding irrelevant failures and exploiting symmetries, see [56] for the details.

Note that modularisation can only be used for computing probabilities, e.g., the unreliability. The MTTF cannot be computed compositionally as combining expectations is difficult [9]. Thus, other approaches are necessary such as the approximation from Sect. 3.3 or composing independent Markov models [56].

Algorithm. We shortly describe our implementation of the DFT analysis via modularisation based on [31]. Algorithm 1 presents the pseudo-code. We use the DFT in Fig. 1 as an example and compute the unreliability within time bound t . We start the analysis by identifying the modules in the DFT using the algorithm of [24]. The algorithm traverses the fault tree in a depth-first left-most order and stores the order in which nodes are visited. A node v is a root of a module if all its descendants are visited in-between the first and last visit of v .

The algorithm runs in linear time and yields a unique list of modules. Minor adaptations of the algorithm are required to adequately handle SEQ and FDEP, cf. [6]. Next, we only keep the dynamic modules, i.e., modules containing at least one dynamic element. We iteratively remove a module if its corresponding elements not contained in other modules only contain static gates. That way, we remove modules containing only static elements and modules which are a subset of dynamic modules. This step results in a unique set of dynamic modules. The example DFT contains two dynamic modules (indicated by dashed blue boxes on the left DFT in Fig. 1). Next, each dynamic module is translated to a Markov model and analysed according to the given metric [56]. The complete dynamic module is then replaced by a single BE which matches the calculated failure probabilities. In our example, BE H' is chosen such that it has the same probability to fail within time bound t than the whole module of H . In the end, the resulting fault tree (on the right in Fig. 1) contains only static elements. Thus, this SFT can be analysed using the BDD approach presented in Sect. 3.

Static modules could of course also be replaced by corresponding BEs. As building the BDD is efficient, we opt to directly analyse the resulting SFT instead. Specific dynamic structures such as the first child of a PAND or SEQ could also be further modularised following [59]. However, the application of these modularisation rules is very limited and results in semi-Markov chains. This approach is therefore not considered here.

Implementation. We implemented the modularisation in STORM-DFT using the BDD implementation described in Sect. 3.4. A DFT is analysed by translating it into a Markov model as in [56]. While STORM-DFT already supports a modularisation, this top-down approach is only applicable to children of the top-level event. In contrast, the new implementation is applicable to dynamic modules located anywhere in the DFT. Moreover, the Markov models for dynamic modules are cached such that multiple queries can be performed on the same model. This is not possible in the previous implementation which regenerates each model for a new metric. The caching is in particular useful when computing multiple time points and exploiting vectorisation on the resulting SFT.

6 Evaluation of DFT Approach

We evaluate the DFT modularisation and compare with existing approaches.

Configurations. We compare the modularisation using BDDs with two existing approaches within STORM-DFT [56]: the translation to a continuous-time Markov chain (CTMC) and the top-down modularisation.

Benchmarks. We use the following DFT benchmarks:

- 68 DFTs from the FFORT benchmark collection [46].
- 40 DFT obtained by using the SFTs from the Aralia benchmark set and replacing one BE by the DFT `ftpp.1-1` from the FFORT benchmark set.

Table 2. DFT benchmark sizes

Benchmark set	#BEs	#Static gates	#Dyn. gates	#BEs mod.	#Static gates mod.
Adapt. SFT	32–1574	26–1628	3	25–1623	21–1623
Adapt. Railway	194–545	153–487	19–54	22–54	40–168
Adapt. VGS	54–99	31–59	6–20	1–79	0–39
FFORT	6–87	1–50	0–44	1–50	0–21

- 8 DFTs modelling infrastructure in railway station areas [57] and slightly adapted to contain modules.
- 8 DFTs modelling configurations for a vehicle guidance system (VGS) [29] and adapted by removing irrelevant FDEPs.

Table 2 gives statistics on these benchmarks: the minimal and maximal number of BEs, static and dynamic gates in the original DFT as well as the numbers for the SFT after modularisation.

We run the three configurations of STORM-DFT on the 124 DFTs. We compute the *unreliability* at a time bound t either given by the largest bound specified in FFORT or we use $t = 100$ otherwise. For *multiple time bounds*, we use 1000 time bounds uniformly distributed over the interval $[0, t]$. We used the same machine and settings (timeout 5 min, 30 GB memory) as in Sect. 4.1.

6.1 Results

Unreliability. Figure 8 compares the computation of the unreliability. Figure 8(a) compares the modularisation using BDDs with the CTMC approach. The new BDD approach solves nearly all DFTs within 1 s and outperforms the CTMC approach by several orders of magnitude. Modularisation is therefore offering clear performance benefits compared to plain CTMC analysis.

Figure 8(b) compares the new BDD-based modularisation with the existing top-down modularisation. The new approach prevails for all larger DFTs. The BDD modularisation solved the adapted SFTs within 0.1 s while the top-down approach required up to 100 s. On these DFTs, top-down modularisation is not applicable and thus, the entire DFT must be translated into a CTMC.

The advantage of the BDD modularisation becomes even clearer for multiple time bounds, cf. Fig. 8(c). The BDD modularisation is significantly faster than the top-down modularisation for most of the considered DFTs. On most DFTs, the BDD modularisation is able to compute 1000 time points within 1 s. The main reasons for the performance improvement on multiple time points are the caching of the intermediate Markov models and the use of vectorisation on the resulting BDD, cf. Sect. 3.2.

The BDD-based modularisation is significantly faster than both the plain CTMC approach and the existing top-down modularisation.

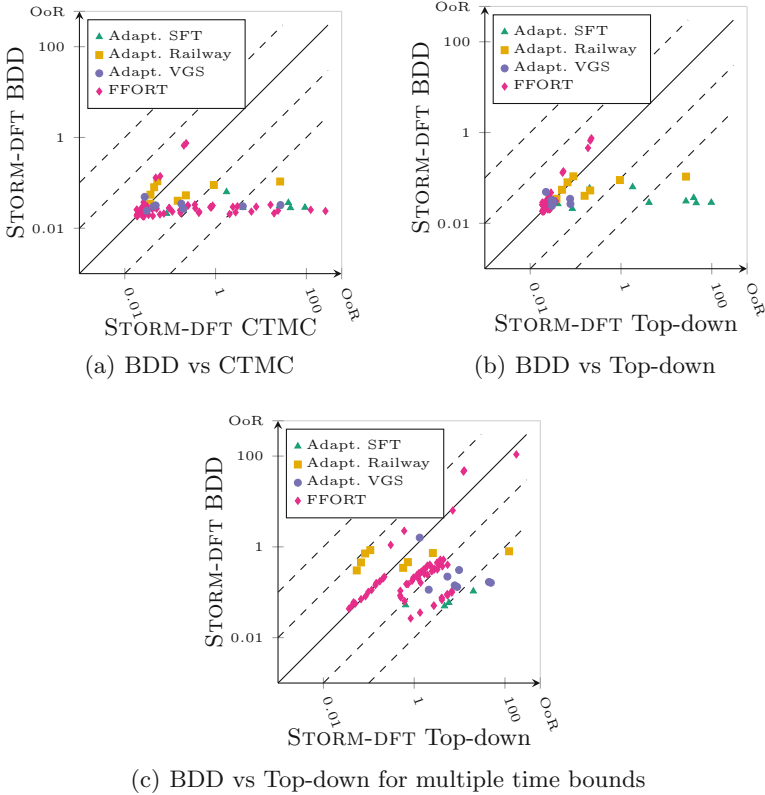


Fig. 8. Comparisons of runtimes for the computation of the unreliability

7 Conclusion

We presented an implementation for fault tree analysis based on BDDs in the STORM-DFT tool. Our implementation is competitive compared to existing tools and performs significantly better when computing multiple time points. We also presented an implementation for DFT analysis based on modularisation. The modular analysis allows to use the best techniques for each sub-tree and outperforms existing approaches. STORM-DFT is currently the only available tool supporting modularisation for efficient DFT analysis.

Future Work. Further improvements are needed to obtain smaller BDDs during the translation, for example by improving the heuristics for variable orderings such as using heuristics from SCRAM. The modularisation cannot be fully exploited if large dynamic modules are present. A possible research direction is to approximate the results for sub-modules, either by smaller fault trees or by the approximation approach of [56].

References

1. Akers, S.B., Jr.: Binary decision diagrams. *IEEE Trans. Comput.* **27**(6), 509–516 (1978)
2. Andrews, J.D., Bartlett, L.M.: Efficient basic event orderings for binary decision diagrams. In: *Annual Reliability and Maintainability Symposium*, pp. 61–68. IEEE (1998)
3. Bäckström, O., Gamble, R., Krcal, P., Wang, W.: An experimental assessment of the MCS BDD algorithm in RiskSpectrum. In: *Safety and Reliability-Safe Societies in a Changing World*, pp. 1709–1717. CRC Press (2018)
4. Bäckström, O., Butkova, Y., Hermanns, H., Krčál, J., Krčál, P.: Effective static and dynamic fault tree analysis. In: Skavhaug, A., Guiochet, J., Bitsch, F. (eds.) *SAFECOMP 2016. LNCS*, vol. 9922, pp. 266–280. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45477-1_21
5. Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.P.: Model-checking algorithms for continuous-time Markov chains. *IEEE Trans. Softw. Eng.* **29**(6), 524–541 (2003)
6. Basgöze, D.: Dynamic fault tree analysis using binary decision diagrams. Bachelor thesis, RWTH Aachen University (2020)
7. Basgöze, D., Volk, M., Katoen, J., Khan, S., Stoelinga, M.: BDDs strike back: efficient analysis of static and dynamic fault trees. *CoRR abs/2202.02829* (2022)
8. Birnbaum, Z.: On the importance of different components in a multicomponent system. *Multivariate Analysis-II*, pp. 581–592 (1969)
9. Bohnenkamp, H., Haverkort, B.: The mean value of the maximum. In: Hermanns, H., Segala, R. (eds.) *PAPM-PROBMIV 2002. LNCS*, vol. 2399, pp. 37–56. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45605-8_4
10. Boudali, H., Crouzen, P., Stoelinga, M.: A rigorous, compositional, and extensible framework for dynamic fault tree analysis. *IEEE Trans. Dependable Secur. Comput.* **7**(2), 128–143 (2010)
11. Bouissou, M., Bruyere, F., Rauzy, A.: BDD based fault-tree processing: a comparison of variable ordering heuristics. In: *European Safety and Reliability Association Conference, ESREL* (1997)
12. Brace, K.S., Rudell, R.L., Bryant, R.E.: Efficient implementation of a BDD package. In: *DAC*, pp. 40–45. IEEE Computer Society Press (1990)
13. Bryant, R.E.: Graph-based algorithms for Boolean function manipulation. *IEEE Trans. Comput.* **35**(8), 677–691 (1986)
14. Bryant, R.E.: Binary decision diagrams. In: *Handbook of Model Checking*, pp. 191–217. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-10575-8_7
15. Budde, C.E., Ruijters, E., Stoelinga, M.: The dynamic fault tree rare event simulator. In: Gribaudo, M., Jansen, D.N., Remke, A. (eds.) *QEST 2020. LNCS*, vol. 12289, pp. 233–238. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59854-9_17
16. Cheok, M.C., Parry, G.W., Sherry, R.R.: Use of importance measures in risk-informed regulatory applications. *Reliab. Eng. Syst. Saf.* **60**(3), 213–226 (1998)
17. Coudert, O., Madre, J.C.: Fault tree analysis: 10/sup 20/prime implicants and beyond. In: *Annual Reliability and Maintainability Symposium*, pp. 240–245. IEEE (1993)
18. Davis, P.J., Rabinowitz, P.: *Methods of Numerical Integration*. Academic Press, Cambridge (1984)
19. van Dijk, T.: *Sylvan: multi-core decision diagrams*. Ph.D. thesis, University of Twente, The Netherlands (2016)

20. van Dijk, T., van de Pol, J.: Sylvan: multi-core framework for decision diagrams. *Int. J. Softw. Tools Technol. Transf.* **19**(6), 675–696 (2017)
21. Dugan, J.B., Bavuso, S.J., Boyd, M.A.: Fault trees and sequence dependencies. In: *Annual Reliability and Maintainability Symposium*, pp. 286–293 (1990)
22. Dugan, J.B., Bavuso, S.J., Boyd, M.A.: Dynamic fault-tree models for fault-tolerant computer systems. *IEEE Trans. Reliab.* **41**(3), 363–377 (1992)
23. Dugan, J.B., Venkataraman, B., Gulati, R.: DIFtree: a software package for the analysis of dynamic fault tree models. In: *Annual Reliability and Maintainability Symposium*, pp. 64–70. IEEE (1997)
24. Dutuit, Y., Rauzy, A.: A linear-time algorithm to find modules of fault trees. *IEEE Trans. Reliab.* **45**(3), 422–425 (1996)
25. Dutuit, Y., Rauzy, A.: Efficient algorithms to assess component and gate importance in fault tree analysis. *Reliab. Eng. Syst. Saf.* **72**(2), 213–222 (2001)
26. Eijkhout, V.: *Introduction to High Performance Scientific Computing*. lulu.com (2011)
27. Federal Aviation Administration: *System safety handbook* (2000)
28. Fussell, J.: How to hand-calculate system reliability and safety characteristics. *IEEE Trans. Reliab.* **24**(3), 169–174 (1975)
29. Ghadhab, M., Junges, S., Katoen, J.P., Kuntz, M., Volk, M.: Safety analysis for vehicle guidance systems with dynamic fault trees. *Reliab. Eng. Syst. Saf.* **186**, 37–50 (2019)
30. Guennebaud, G., Jacob, B., et al.: *Eigen v3* (2010). <http://eigen.tuxfamily.org>
31. Gulati, R., Dugan, J.B.: A modular approach for analyzing static and dynamic fault trees. In: *Annual Reliability and Maintainability Symposium*, pp. 57–63. IEEE (1997)
32. Hensel, C., Junges, S., Katoen, J.P., Quatmann, T., Volk, M.: The probabilistic model checker storm. *Int. J. Softw. Tools Technol. Transf.* 1–22 (2021). <https://doi.org/10.1007/s10009-021-00633-z>
33. ISO: *ISO 26262: Road vehicles - Functional safety*. Standard, International Organization for Standardization, Geneva, Switzerland (2011)
34. Junges, S., Katoen, J.-P., Stoelinga, M., Volk, M.: One net fits all. In: Khomenko, V., Roux, O.H. (eds.) *PETRI NETS 2018*. LNCS, vol. 10877, pp. 272–293. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-91268-4_14
35. Krcál, J., Krcál, P.: Scalable analysis of fault trees with dynamic features. In: *DSN*, pp. 89–100. IEEE Computer Society (2015)
36. Mo, Y.: A multiple-valued decision-diagram-based approach to solve dynamic fault trees. *IEEE Trans. Reliab.* **63**(1), 81–93 (2014)
37. Moinuddin, K., Innocent, J., Keshavarz, K.: Reliability of sprinkler system in Australian shopping centres—a fault tree analysis. *Fire Saf. J.* **105**, 204–215 (2019)
38. Montani, S., Portinale, L., Bobbio, A., Raiteri, D.C.: Radyban: a tool for reliability analysis of dynamic fault trees through conversion into dynamic Bayesian networks. *Reliab. Eng. Syst. Saf.* **93**(7), 922–932 (2008)
39. Raiteri, D.C.: The conversion of dynamic fault trees to stochastic Petri nets, as a case of graph transformation. *Electron. Notes Theor. Comput. Sci.* **127**(2), 45–60 (2005)
40. Rakhimov, O.: *Scram probabilistic risk analysis tool* (2018). <https://doi.org/10.5281/zenodo.1146337>
41. Rao, K.D., Gopika, V., Rao, V.V.S.S., Kushwaha, H.S., Verma, A.K., Srividya, A.: Dynamic fault tree analysis using Monte Carlo simulation in probabilistic safety assessment. *Reliab. Eng. Syst. Saf.* **94**(4), 872–883 (2009)

42. Rauzy, A.: Some disturbing facts about depth-first left-most variable ordering heuristics for binary decision diagrams. *Proc. Inst. Mech. Eng. Part O: J. Risk Reliab.* **222**(4), 573–582 (2008)
43. Rauzy, A.: New algorithms for fault trees analysis. *Reliab. Eng. Syst. Saf.* **40**(3), 203–211 (1993)
44. Rauzy, A.: Sequence algebra, sequence decision diagrams and dynamic fault trees. *Reliab. Eng. Syst. Saf.* **96**(7), 785–792 (2011)
45. Rauzy, A.: Probabilistic safety analysis with XFTA. AltaRica Association (2020)
46. Ruijters, E., et al.: FFORT: a benchmark suite for fault tree analysis. In: ESREL. Singapore: Research Publishing (2019)
47. Ruijters, E., Stoelinga, M.: Fault tree analysis: a survey of the state-of-the-art in modeling, analysis and tools. *Comput. Sci. Rev.* **15**, 29–62 (2015)
48. Shannon, C.E.: A symbolic analysis of relay and switching circuits. *Electr. Eng.* **57**(12), 713–723 (1938)
49. Sinnamon, R.M., Andrews, J.: Improved efficiency in qualitative fault tree analysis. *Qual. Reliab. Eng. Int.* **13**(5), 293–298 (1997)
50. Sinnamon, R.M., Andrews, J.D.: Fault tree analysis and binary decision diagrams. In: Annual Reliability and Maintainability Symposium, pp. 215–222. IEEE (1996)
51. Stamatelatos, M., Vesely, W., Dugan, J., Fragola, J., Minarick, J., Railsback, J.: *Fault Tree Handbook with Aerospace Applications*. NASA, Washington, DC (2002)
52. Steven, E., Antoine, R.: Open-PSA Model Exchange Format. The Open-PSA Initiative (2007)
53. Sullivan, K.J., Dugan, J.B., Coppit, D.: The Galileo fault tree analysis tool. In: FTCS, pp. 232–235. IEEE Computer Society (1999)
54. Trivedi, K.S., Bobbio, A.: *Reliability and Availability Engineering - Modeling, Analysis, and Applications*. Cambridge University Press, Cambridge (2017)
55. Trivedi, K.S., Sahner, R.A.: SHARPE at the age of twenty two. *SIGMETRICS Perform. Eval. Rev.* **36**(4), 52–57 (2009)
56. Volk, M., Junges, S., Katoen, J.P.: Fast dynamic fault tree analysis by model checking techniques. *IEEE Trans. Ind. Inform.* **14**(1), 370–379 (2018)
57. Volk, M., Weik, N., Katoen, J.-P., Nießen, N.: A DFT modeling approach for infrastructure reliability analysis of railway station areas. In: Larsen, K.G., Willemse, T. (eds.) FMICS 2019. LNCS, vol. 11687, pp. 40–58. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-27008-7_3
58. Xing, L., Tannous, O., Dugan, J.B.: Reliability analysis of nonrepairable cold-standby systems using sequential binary decision diagrams. *IEEE Trans. Syst. Man Cybern. Part A* **42**(3), 715–726 (2012)
59. Yevkin, O.: An improved modular approach for dynamic fault tree analysis. In: Annual Reliability and Maintainability Symposium, pp. 1–5. IEEE (2011)
60. Zhou, S., Xiang, J., Wong, W.E.: Reliability analysis of dynamic fault trees with spare gates using conditional binary decision diagrams. *J. Syst. Softw.* **170**, 110766 (2020)